AMENDMENTS TO THE CLAIMS

Claims 1-11 (cancelled)

Claim 12 (new): A method for mapping a virtual address space into block addresses of at least one data storage device, the method comprising:

    generating a hierarchical data structure in a primary storage;

    wherein the hierarchical data structure includes a plurality of layers arranged according to a hierarchy;

    wherein the plurality of layers include at least a highest layer and a lowest layer;

    wherein each layer in the hierarchical data structure includes at least one set of data entries;

    wherein each data entry in each layer represents a range of the virtual address space;

    wherein for each layer in the hierarchical data structure for which there exists an next lowest layer, each data entry is correlated to a set of data entries in the next lowest layer according to a correlation scheme;

    wherein each data entry in the lowest layer corresponds to both a virtual address range in the virtual address space and a block address corresponding to a physical data block in the at least one data storage device;

    wherein each data entry contained within the primary storage corresponds to a virtual address range that is currently occupied with stored data, such that none of the data entries corresponds to only unused physical storage; and

    wherein each physical data block in the at least one data storage device contains virtual address information that identifies at least one corresponding location in the virtual address space for that physical data block.

Claim 13 (new): The method of claim 12, further comprising:

    swapping sub-hierarchies of data entries between the hierarchical data structure in primary storage and a secondary storage.

Claim 14 (new): The method of claim 12, wherein at least some of the data entries in the lowest layer include a device address identifying an individual storage device in the at least one data storage device.

Claim 15 (new): The method of claim 12, wherein the correlation scheme is one of an algorithm, a hash algorithm, a pointer system, and a pointer to correlation logic.

Claim 16 (new): The method of claim 12, wherein at least some of the data entries in each layer represent virtual address ranges of a homogeneous size corresponding to that layer.

Claim 17 (new): The method of claim 16, further comprising:

generating a second data structure, wherein the second data structure identifies exceptional data entries in the hierarchical data structure, wherein each individual exception data entry corresponds to a virtual address range a size that differs from the homogeneous size corresponding to that layer to which the individual exceptional data entry belongs.

Claim 18 (new): The method of claim 16, wherein the virtual address information identifies a plurality of corresponding locations in the virtual address space for the physical data block.

Claim 19 (new): A computer program product in a computer-readable medium for mapping a virtual address space into block addresses of at least one data storage device, the computer program product comprising:

instructions for generating a hierarchical data structure in a primary storage;

wherein the hierarchical data structure includes a plurality of layers arranged according to a hierarchy;

wherein the plurality of layers include at least a highest layer and a lowest layer;

wherein each layer in the hierarchical data structure includes at least one set of data entries;

wherein each data entry in each layer represents a range of the virtual address space;

wherein for each layer in the hierarchical data structure for which there exists an next lowest layer, each data entry is correlated to a set of data entries in the next lowest layer according to a correlation scheme;

wherein each data entry in the lowest layer corresponds to both a virtual address range in the virtual address space and a block address corresponding to a physical data block in the at least one data storage device;

wherein each data entry contained within the primary storage corresponds to a virtual address range that is currently occupied with stored data, such that none of the data entries corresponds to only unused physical storage; and

wherein each physical data block in the at least one data storage device contains virtual address information that identifies at least one corresponding location in the virtual address space for that physical data block.

Claim 20 (new): The computer program product of claim 19, further comprising:

instructions for swapping sub-hierarchies of data entries between the hierarchical data structure in primary storage and a secondary storage.

Claim 21 (new): The computer program product of claim 19, wherein at least some of the data entries in the lowest layer include a device address identifying an individual storage device in the at least one data storage device.

Claim 22 (new): The computer program product of claim 19, wherein the correlation scheme is one of an algorithm, a hash algorithm, a pointer system, and a pointer to correlation logic.

Claim 23 (new): The computer program product of claim 19, wherein at least some of the data entries in each layer represent virtual address ranges of a homogeneous size corresponding to that layer.

Claim 24 (new): The computer program product of claim 23, further comprising:

instructions for generating a second data structure, wherein the second data structure identifies exceptional data entries in the hierarchical data structure, wherein each individual exception data entry corresponds to a virtual address range a size that differs from the homogeneous size corresponding to that layer to which the individual exceptional data entry belongs.

Claim 25 (new): The computer program product of claim 19, wherein the virtual address information identifies a plurality of corresponding locations in the virtual address space for the physical data block.

Claim 26 (new): A data management system for mapping a virtual address space into block addresses of at least one data storage device, the data management system comprising:

means for generating a hierarchical data structure in a primary storage;

wherein the hierarchical data structure includes a plurality of layers arranged according to a hierarchy;

wherein the plurality of layers include at least a highest layer and a lowest layer;

wherein each layer in the hierarchical data structure includes at least one set of data entries;

wherein each data entry in each layer represents a range of the virtual address space;

wherein for each layer in the hierarchical data structure for which there exists an next lowest layer, each data entry is correlated to a set of data entries in the next lowest layer according to a correlation scheme;

wherein each data entry in the lowest layer corresponds to both a virtual address range in the virtual address space and a block address corresponding to a physical data block in the at least one data storage device;

wherein each data entry contained within the primary storage corresponds to a virtual address range that is currently occupied with stored data, such that none of the data entries corresponds to only unused physical storage; and

wherein each physical data block in the at least one data storage device contains virtual address information that identifies at least one corresponding location in the virtual address space for that physical data block.

Claim 27 (new): The data management system of claim 26, further comprising:

means for swapping sub-hierarchies of data entries between the hierarchical data structure in primary storage and a secondary storage.

Claim 28 (new): The data management system of claim 26, wherein at least some of the data entries in the lowest layer include a device address identifying an individual storage device in the at least one data storage device.

Claim 29 (new): The data management system of claim 26, wherein at least some of the data entries in each layer represent virtual address ranges of a homogeneous size corresponding to that layer.

Claim 30 (new): The data management system of claim 29, further comprising:

means for generating a second data structure, wherein the second data structure identifies exceptional data entries in the hierarchical data structure, wherein each individual exception data entry corresponds to a virtual address range a size that differs from the homogeneous size corresponding to that layer to which the individual exceptional data entry belongs.

Claim 31 (new): The data management system of claim 26, wherein the virtual address information identifies a plurality of corresponding locations in the virtual address space for the physical data block.